

Bounds on the Rate of Linear Locally Repairable Codes over Small Alphabets

Abhishek Agarwal and Arya Mazumdar *

Abstract

Locally repairable codes (LRC) have recently been a subject of intense research due to theoretical appeal and their application in distributed storage systems. In an LRC, any coordinate of a codeword can be recovered by accessing only few other coordinates. For LRCs over small alphabet (such as binary), the optimal rate-distance trade-off is unknown. In this paper we provide the tightest known upper bound on the rate of linear LRCs of a given relative distance, an improvement over any previous result, in particular [1].

1 Introduction

Let $\mathcal{C}(n, d) \subseteq \{0, 1\}^n$ denote a binary linear code of length n with minimum distance $d_{\min} = d$ and \mathcal{C}^\perp denote the dual of the linear code \mathcal{C} . A code $\mathcal{C}(n, d)$ can recover any codeword after the erasure of at most $(d - 1)$ coordinate symbol. But, it requires knowledge of all other $n - d + 1$ coordinates of the codeword for recovery. For certain scenarios such as distributed storage systems, in addition to minimum distance, codes with a small *locality* are desired. In such systems, to repair a single symbol of a codeword of \mathcal{C} , we want to utilize information about few other codeword symbols. Locally Repairable Codes (LRC) [5], with locality parameter r , allow the repair of any one symbol of a codeword by knowing at most r other codeword symbols. Let S be the subset of coordinates (of size at most r) required to repair a symbol at coordinate i in the LRC. Then the set $R(i) := S \cup \{i\}$ is called the repair group of coordinate i . We denote the linear LRC of length n , minimum distance d and locality r as $\mathcal{C}(n, d, r)$.

The dimension of $\mathcal{C}(n, d)$ is defined as $k(\mathcal{C}) = \log_2 |\mathcal{C}|$, and the rate as $R(\mathcal{C}) = \frac{k}{n}$. For a code $\mathcal{C}(n, d, r)$, its dimension is known to be upper bounded by,

$$k(\mathcal{C}) \leq m - \left\lfloor \frac{m}{r+1} \right\rfloor \quad (1)$$

where $m = n - d + 1$ [5]. This bound is known to be achievable when the code is over alphabet size at least n [9], but it is far from being tight for smaller, in particular binary, alphabets [1]. For binary code $\mathcal{C}(n, d, r)$,

*Abhishek Agarwal is with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455. Arya Mazumdar is with the Computer Science Department of University of Massachusetts, Amherst, MA 01003, and was with the University of Minnesota. email: abhiag@umn.edu, arya@cs.umass.edu. Research supported by NSF grants CCF 1318093, CCF 1642658, CCF 1453121 and CCF 1618512.

an upper bound on the dimension of LRC was presented in [1]:

$$k \leq \min_{t \in \mathbb{Z}_+} \left[tr + k_{\text{opt}}(n - t(r + 1), d) \right], \quad (2)$$

where $k_{\text{opt}}(n, d)$ is the optimum dimension of an error-correcting code of length n and distance d .

This bound can be translated into an asymptotic bound on rate and relative distance, assuming $k/n \rightarrow R$ and $d/n \rightarrow \delta$ as $n \rightarrow \infty$:

$$R \leq \min_{0 \leq x \leq r/(r+1)} x + (1 - x(1 + 1/r)) \cdot R_{\text{opt}} \left(\frac{\delta}{1 - x(1 + 1/r)} \right) \quad (3)$$

where, $R_{\text{opt}}(\delta) \equiv \lim_{n \rightarrow \infty} \frac{k_{\text{opt}}(n, \delta n)}{n}$. The best known upper bound on the rate of a locally repairable code is found when $R_{\text{opt}}(\delta)$ is bounded by $R_{\text{MRRW}}(\delta)$ given below:

$$R_{\text{MRRW}}(\delta) = \min_{0 < \alpha \leq 1 - 4\delta} 1 + \hat{h}(\alpha^2) - \hat{h}(\alpha^2 + 4\delta\alpha + 4\delta), \quad (4)$$

with $\hat{h}(x) = h(1/2 - 1/2\sqrt{1-x})$ and $h(x) \equiv -x \log_2 x - (1-x) \log_2(1-x)$ is the binary entropy function.

A Gilbert-Varshamov type achievability bound was also presented in [2]. These best known upper and lower bounds are well-separated. Despite intense research on LRCs, and in particular on rate-bounds (cf. [6]), no result asymptotically better than (3) is known. In this paper, we provide improved upper bounds on the rate of linear LRCs that are tighter than (3), for a range of rate-relative distance trade-off.

Our algebraic combinatorics methods are inspired by a recent work [7], that provides an improved estimate of rate of a code with sparse parity-check matrix. A combination of techniques from [1] and [7] leads to improvement over the bounds in [1] for linear LRCs with arbitrary repair groups of size at most r . We further improve over these bounds for disjoint repair groups. Our bounds apply for the asymptotic case and give an improvement over the best known bounds on linear LRC for small alphabets. For clarity, we state our result for binary alphabet.

The paper is organized as follows. In Section 2.1, we provide the main theorems: bounds on the rate of linear LRCs in general, and for the special case of disjoint repair groups. The main techniques of coset counting is presented in Section 2.2. The proof of the main results are presented subsequently.

Notations

For integers a, b the relation $a \mid b$ denotes that a divides b . $[m, n] := \{m, m+1, \dots, n\}$ and $[n] := \{1, 2, \dots, n\}$. $\sqcup_i A_i$ denotes the union of disjoint sets A_i . $H((p_1, p_2, \dots, p_n)^T) := \sum_i -p_i \log p_i$ denotes the entropy for the distribution $\{p_i\}_i$ while $h(p) := -p \log p - (1-p) \log(1-p)$ denotes the binary entropy. The logarithms $\log(\cdot)$ are taken to be base 2 throughout the paper. For a vector $\mathbf{v} \in \mathbb{F}_2^n$, let $\mathbf{v}(i)$ denote the symbol at coordinate i of \mathbf{v} . Denote $s(\mathbf{v})$ as the set $\{i \in [n] \mid \mathbf{v}(i) \neq 0\}$ and $\text{wt}(\mathbf{v}) := |s(\mathbf{v})|$. Define $\mathbf{e}_i, i \in [n]$ as a vector in \mathbb{F}_2^n such that $s(\mathbf{v}) = \{i\}$.

2 Rate upper bounds for LRC

2.1 Main results

Our first result for general linear LRC rate is presented below.

Theorem 1 (Linear LRC). *The rate of a binary linear LRC $\mathcal{C}(n, \delta n, r)$ can be at most,*

$$R(\mathcal{C}) \leq \min_{0 \leq x \leq r/(r+1)} x + (1 - x(1 + 1/r)) \cdot R_1\left(\frac{\delta}{1 - x(1 + 1/r)}, r\right),$$

where, $R_1(\delta, r) = h(\rho) - c(r + 1, \rho)$, and $c(w, \rho) := \frac{\log_2 e}{8w^2} \left(\frac{\rho^w}{2}\right)^{w+1}$.

This bound is an improvement on (3) with R_{opt} replaced by R_{MRRW} and R_1 respectively. For LRCs with disjoint repair groups we have an even tighter bound on the rate.

Theorem 2 (Linear LRC with Disjoint Repair Groups). *Let,*

$$\beta_i(x) := \begin{cases} \frac{\binom{r+1}{i} x^i}{\binom{r+1}{i}} & i < t \\ \frac{\binom{r+1}{i}}{2^{r \bmod 2} x^i} & i = t \end{cases} \quad (5)$$

The asymptotic rate $R(\mathcal{C})$ for a binary linear LRC $\mathcal{C}(n, \delta n, r)$ with disjoint repair groups satisfies,

$$R(\mathcal{C}) \leq \rho \log(x) + \frac{\log(\sum_{i=0}^t \beta_i(x))}{r+1} \Big|_{x=\mu} \quad (6)$$

where $\mu = 1$ when $\sum_i i \frac{\beta_i(x)}{\sum_i \beta_i(x)} \Big|_{x=1} \leq (r+1)\rho$ and

$$\mu = \text{Roots}^+((r+1)\rho \sum_{i=0}^t \binom{r+1}{i} x^{t-i} - \sum_{i=0}^t \binom{r+1}{i} i x^{t-i}) \quad (7)$$

otherwise and where $t := \lfloor \frac{r+1}{2} \rfloor$ and $\text{Roots}^+(f(x))$ denotes the unique positive root of $f(x) = 0$.

In fig. 1 we compare the result of theorem 2 with the existing asymptotic bounds. The bounds converge together as δ and r increase. But we can see from fig. 2 that our bound improves over the bound in (3) (cf. [1]) for $\delta \geq 0.38$ for $r = 2$ and this range improves as r increases.

2.2 The method of counting cosets

The first MRRW bound provides the tightest (except for the second MRRW bound) upper bound on the rate of a code $\mathcal{C}(n, \delta)$:

$$R(\mathcal{C}) \leq h(\rho), \quad (8)$$

where $\rho := \frac{1}{2} - \sqrt{\delta(1-\delta)}$. The bound is derived in [8] using linear constraints on the weight (distance) distribution of the code \mathcal{C} . Another way to derive the first MRRW bound for linear codes is to analyze the distance between the codewords in the *coset leader graph* of \mathcal{C} [4]. These graphs are defined as follows [4].




Figure 1: Comparison of the asymptotic upper-bound on rate (vs relative minimum distance) in eq. (6) with the existing bounds in [1, 5, 8].




Figure 2: Difference between the bound in (3) (cf. [1]) and our bound eq. (6).

Definition 1. The coset leader graph $\mathbb{T}(\mathcal{C}) = (\mathcal{V}, \mathcal{E})$ of a binary linear code \mathcal{C} is a graph with vertex set, $\mathcal{V} = \{\mathbf{x} + \mathcal{C}^\perp \mid \mathbf{x} \in \mathbb{F}_2^n\}$ as the cosets of \mathcal{C}^\perp . Two vertices in $\mathbb{T}(\mathcal{C})$ are connected iff the corresponding cosets are Hamming distance 1 apart i.e. $\{\mathbf{x} + \mathcal{C}^\perp, \mathbf{y} + \mathcal{C}^\perp\} \in \mathcal{E} \implies \mathbf{x} = \mathbf{y} + \mathbf{e}_i + \mathbf{v}$, for some $\mathbf{v} \in \mathcal{C}^\perp$ and $i \in [n]$.

It is known [3] that the eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots$ of the adjacency matrix of $\mathbb{T}(\mathcal{C})$ satisfy,

$$\lambda_i = n - 2d_i$$

where $d_1 \leq d_2 \leq \dots$ are the weights of the codewords in \mathcal{C} . Since $d_1 = 0$ for any linear code, we have $\delta = \frac{1}{2}(1 - \lambda_2/n)$, i.e., the minimum distance of the code corresponds to the second eigenvalue of the coset leader graph.

Note that, the shortest path length between the cosets \mathcal{C}^\perp and $\mathbf{x} + \mathcal{C}^\perp$ is equal to the weight of the coset leader of $\mathbf{x} + \mathcal{C}^\perp$. Denote the distance between two vertices in the graph $\mathbb{T}(\mathcal{C})$ as the length of the shortest path between them. Let $B_{\mathbb{T}}(v, r)$ denote the vertices in $\mathbb{T}(\mathcal{C})$ inside a ball of radius r and centered at vertex v . Since, the graph $\mathbb{T}(\mathcal{C})$ is vertex transitive, the volume in the $B_{\mathbb{T}}(v, r)$ does not depend on v and is written as $B_{\mathbb{T}}(r)$. Thus, $B_{\mathbb{T}}(r)$ also denotes the number of coset leaders of \mathcal{C}^\perp of weight not greater than r . From [7, Theorem 1.1], we know that the asymptotic rate, $R(\mathcal{C})$ of code \mathcal{C} with relative minimum distance δ , can be upper-bounded using the coset leader graph $\mathbb{T}(\mathcal{C})$ as,

$$|\mathcal{C}| \leq 2^{o(n)} |B_{\mathbb{T}} \rho n|, \quad (9)$$

where, recall, $\rho := \frac{1}{2} - \sqrt{\delta(1 - \delta)}$. Thus, the asymptotic MRRW bound on $R(\mathcal{C})$ in eq. (8) can be derived using a very general bound,

$$B_{\mathbb{T}} \rho n \leq \sum_{i \leq \rho n} \binom{n}{i}. \quad (10)$$

Note that the binary linear LRC $\mathcal{C}(n, \delta n, r)$ also satisfies eq. (9). To improve the rate upper-bound eq. (8) for these codes, we can use their locality property to improve the counting upper-bound on $B_{\mathbb{T}} \rho n$, eq. (10).

2.3 Upper bound for linear LRCs: proof of theorem 1

Consider a binary linear code \mathcal{C} with parity check matrix of constant row weights w . Then, for this code \mathcal{C} the following bound holds [7, Thm 1.2],

$$|B_{\mathbb{T}} \rho n| \leq 2^{-c(w, \rho)n} |B(\rho n)| \quad (11)$$

$B(t)$ denotes a ball of radius t in \mathbb{F}_2^n . We describe a sketch of the proof of this bound from [7] below. Our proof of theorem 1 will then follow.

Consider a random binary vector \mathbf{v} such that $\mathbf{v}(i) \sim \text{Bernoulli}(\rho)$ are iid for $i \in [n]$. Then the probability $p(\rho)$ that the vector \mathbf{v} is a coset leader of \mathcal{C}^\perp satisfies [7, Lemma 2.1],

$$p(\rho) \geq \Omega\left(\frac{1}{\sqrt{n}}\right) \frac{|B_{\mathbb{T}} \rho n|}{|B(\rho n)|}. \quad (12)$$

We now upper bound the probability $p(\rho)$ by determining a necessary condition for each coset leader. Next, we use eq. (12) and the upper bound on $p(\rho)$ to find an upper bound on $B_{\mathbb{T}} \rho n$.

Let the parity check matrix for the code \mathcal{C} be $H = [\mathbf{v}_1 \mathbf{v}_2 \dots \mathbf{v}_m]^T$ such that $\text{wt}(\mathbf{v}_i) = w, \forall i$. Then we divide the coordinates $[n]$ into partitions $\{I_j\}_{j \in [w]}$ using a recursive algorithm. Assume that $I_l, l > j$ have already been constructed. Then I_j is constructed as in algorithm 1. Note that each partition I_j is such that $I_j = \bigsqcup_{l=1}^t U_{jl}$ for $U_{jl} : |U_{jl}| = j$. Let the $\mathbf{v}_{jl} \in \{\mathbf{v}_i\}_i$ be the vector containing the support of the U_{ji} . Then, it is easy to see that $s(\mathbf{v}_{ji}) \setminus U_{ji} \subseteq \bigcup_{i=1}^{j+1} I_i$.

Data: $\{\mathbf{v}_i\}_{i \in [m]}, \{I_l\}_{l=j+1}^w$
Result: I_j

```

1  $I_j = \emptyset, T = \{\mathbf{v}_i : i \in [m]\};$ 
2 while  $\exists \mathbf{v}_i \in T$  do
3   if  $|s(\mathbf{v}_i) \setminus \bigcup_{l=j}^w I_l| = j$  then
4      $I_j = I_j \cup \{s(\mathbf{v}_i) \setminus \bigcup_{l=j}^w I_l\};$ 
5   end
6    $T = T \setminus \{\mathbf{v}_i\}$ 
7 end
```

Algorithm 1: Recursive algorithm for partitioning the coordinates in the code \mathcal{C}

An important property of the partitions $\{I_j\}_j$ is as follows [7, Lemma 2.2],

Lemma 3. *There exists a partition I_k such that,*

$$|I_k| \geq \max \left\{ A \sum_{j>k} |I_j|, B \right\}$$

where $A = \frac{2}{\rho^w}$ and $B = \frac{n}{2wA^w}$.

The proof of lemma 3 follows from the requirement that there must exist at least one $k \in [n]$ such that $|I_k| \geq n/w$. Let the coordinate k be such that lemma 3 is satisfied i.e. $|I_k| \geq \max \left\{ A \sum_{j>k} |I_j|, B \right\}$. Then, we claim that the support for every coset leader \mathbf{x} can contain at most $\frac{\rho^k}{2}t$ of the k -tuples $\{U_{kl}\}_{l=1}^t$. This can be immediately seen, since

$$\mathbf{y} = \mathbf{x} + \sum_{i \in S} \mathbf{v}_{ki}$$

has weight less than \mathbf{x} for the set $S := \{i \mid U_{ki} \subseteq s(\mathbf{x})\}$ such that $|S| \geq \frac{\rho^k}{2}$. Thus, using this necessary condition for the coset leader \mathbf{x} we can upper bound $p(\rho)$ in eq. (12) as follows,

$$\begin{aligned}
p(\rho) &\leq \mathbb{P} \left(\text{Binomial}(\rho^k, t) \leq \frac{\rho^k}{2}t \right) \leq \exp \left(-\frac{\rho^k \cdot t}{8} \right) \\
&\stackrel{(a)}{\leq} 2^{-c(w, \rho)n},
\end{aligned} \tag{13}$$

where inequality (a) above follows since for k satisfying lemma 3 we have,

$$\rho^k \cdot t \leq \frac{\rho^w}{w} B = \frac{\rho^w}{w} \cdot \frac{1}{2wA^w} \cdot n.$$

Here, we note that the upper-bound eq. (11) requires only the assumption that the parity check matrix H of \mathcal{C} contains constant weight row vectors \mathbf{v}_i whose support covers the coordinates $[n]$. Note that for an LRC, every coordinate must have a repair group of size at most $r+1$. Hence in the parity check matrix, there must be rows of weight $r+1$ that covers all the coordinates by their support. Thus, we have the following upper bound for an LRC $\mathcal{C}(n, \delta n, r)$.

Lemma 4. *The asymptotic rate $R(\mathcal{C})$ for a binary linear LRC $\mathcal{C}(n, \delta n, r)$ satisfies,*

$$R(\mathcal{C}) \leq \underbrace{h(\rho) - c(r+1, \rho)}_{R_1(\delta, r)}$$

where and $c(w, \rho) := \frac{\log_2 e}{8w^2} \left(\frac{\rho^w}{2} \right)^{w+1}$.

The upper bound in lemma 4 gives an upper-bound worse than the bound in (3). But we can construct a better bound by using the bound in lemma 4 inside the shortening bound in (3). A mild change in the argument of [1] first leads to the following bound.

Theorem 5. *The rate for an LRC code $\mathcal{C}(n, \delta n, r)$ satisfies,*

$$R(\mathcal{C}) \leq \min_{0 \leq x \leq r/(r+1)} x + (1 - x(1 + 1/r)) \cdot R_{\text{opt}}\left(\frac{\delta}{1 - x(1 + 1/r)}, r\right)$$

where $R_{\text{opt}}(\delta, r)$ is the optimal asymptotic rate of $\mathcal{C}(n, \delta n, r)$.

In [1], the MRRW bound for $R_{\text{opt}}(x, r)$ was used. If we use the bound in lemma 4 for $R_{\text{opt}}(\cdot)$ we have,

$$R(\mathcal{C}) \leq \min_{0 \leq x \leq r/(r+1)} x + (1 - x(1 + 1/r)) \cdot R_1\left(\frac{\delta}{1 - x(1 + 1/r)}, r\right)$$

where $R_1(\cdot)$ is as defined in lemma 4.

This completes the proof of theorem 1.

2.4 Upper bound on LRC for disjoint repair groups: proof of theorem 2

In this section we consider the special case of LRC codes with disjoint repair groups. For this case too, the bound in lemma 4 applies. But we can improve that bound significantly in this situation. Consider a binary linear LRC $\mathcal{C}(n, \delta n, r)$ with disjoint repair groups $R_j, j \in \left[\frac{n}{r+1} \right]$. For clarity, in this section we assume

$(r+1) \mid n$. The parity check-matrix of \mathcal{C} must be of the form,

$$H = \begin{pmatrix} \mathbf{v}_1 \\ \vdots \\ \mathbf{v}_{n'} \\ H' \end{pmatrix}$$

where $n' = \frac{n}{r+1}$, $s(\mathbf{v}_i) \cap s(\mathbf{v}_j) = \emptyset$ and $R_i := s(\mathbf{v}_i)$ are such that $|R_i| = r+1$ (even if the repair groups are of smaller size, we can add redundant coordinates in them to equal $r+1$). Then, $\mathbf{v} \in B_{\mathbb{T}} \rho n$ satisfies the following constraint,

$$\text{wt}(\mathbf{v}) \leq \min_{S \subseteq [n']} \text{wt} \left(\mathbf{v} + \sum_{i \in S} \mathbf{v}_i \right) \quad (14a)$$

$$\text{wt}(\mathbf{v}) \leq \rho n \quad (14b)$$

$$|B_{\mathbb{T}} \rho n \cap (\mathbf{x} + \text{span}(\{\mathbf{v}_i\}_i))| \leq 1 \quad \forall \mathbf{x} \in \mathbb{F}_2^n. \quad (14c)$$

We use eq. (14) to find a better upper bound on $R(\mathcal{C})$.

Proof of theorem 2. We find all vectors satisfying eq. (14) to upper bound $|B_{\mathbb{T}} \rho n|$ and to subsequently get an upper bound on $R(\mathcal{C}) = \log |B_{\mathbb{T}} \rho n|/n$. Using eqs. (14a) and (14b) we must have,

$$B_{\mathbb{T}} \rho n \subseteq \underbrace{\left\{ \mathbf{x} \in \mathbb{F}_2^n \mid \text{wt } \mathbf{x} \leq \rho n, |s(\mathbf{x}) \cap R_i| \leq \left\lfloor \frac{r+1}{2} \right\rfloor \quad \forall R_i \right\}}_{\mathcal{S}}. \quad (15)$$

The number of words in \mathcal{S} is given by

$$\sum_{k \leq \rho n} [x^k] \left(1 + \sum_{i=0}^t \binom{r+1}{i} x^i \right)^{\frac{n}{r+1}} \quad (16)$$

where $t = \lfloor \frac{r+1}{2} \rfloor$ and $[x^k]g(x)$ denotes the coefficient of x^k in $g(x)$.

We can further reduce the bound in eq. (15) using eq. (14c). Since, \mathbf{v} and $\mathbf{v} + \mathbf{v}_i$ are of equal weight and are in the same coset we need to count only one of them in \mathcal{S} . Thus, for $r+1 = \text{even}$, we can improve the bound in eq. (16) as follows,

$$|B_{\mathbb{T}} \rho n| \leq \sum_{k \leq \rho n} [x^k] \underbrace{\left(1 + \sum_{i=0}^{t-1} \binom{r+1}{i} x^i + \frac{\binom{r+1}{t}}{2^{r \bmod 2}} x^t \right)}_{g(x)}^{\frac{n}{r+1}}. \quad (17)$$

We now see that the upper bound in eq. (17) is dominated (asymptotically) by $\max_k [x^k]g(x)$ i.e.

$$\frac{\log \sum_{k \leq \rho n} [x^k]g(x)}{n} = \max_{k \leq \rho n} \frac{\log [x^k]g(x)}{n},$$

as $n \rightarrow \infty$. Thus we have,

$$\begin{aligned}
R(\mathcal{C}) &\leq \max_{k \leq \rho n} \frac{\log [x^k]g(x)}{n} \\
&= \max_{\substack{\sum_i \alpha_i = \frac{1}{r+1} \\ \sum_i i \alpha_i \leq \rho}} \frac{1}{r+1} H\left((r+1) \cdot \begin{pmatrix} \alpha_0 \\ \vdots \\ \alpha_t \end{pmatrix}\right) \\
&\quad + \sum_{i=0}^{t-1} \alpha_i \log \binom{r+1}{i} + \alpha_t \log \frac{\binom{r+1}{t}}{2^{r \bmod 2}}.
\end{aligned} \tag{18}$$

Since, eq. (18) represents a concave optimization with linear constraints, we can find the optima in eq. (18) analytically. The optimum value is achieved at $\alpha_i = \frac{\beta_i(x)}{(r+1) \sum_{i=1}^t \beta_i(x)} \big|_{x=\mu}$ for $\beta_i(x)$ and μ as defined in eqs. (5) and (7) and eq. (18) evaluates to eq. (6). \square

References

- [1] V. Cadambe and A. Mazumdar. An upper bound on the size of locally recoverable codes. In *Proc. IEEE Int. Symp. Network Coding*, June 2013.
- [2] V. Cadambe and A. Mazumdar. Bounds on the size of locally recoverable codes. *Information Theory, IEEE Transactions on*, 61(11):5787–5794, 2015.
- [3] C. Delorme and P. Solé. Diameter, covering index, covering radius and eigenvalues. *European Journal of Combinatorics*, 12(2):95–108, 1991.
- [4] J. Friedman and J.-P. Tillich. Generalized alon–boppana theorems and error-correcting codes. *SIAM Journal on Discrete Mathematics*, 19(3):700–718, 2005.
- [5] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin. On the locality of codeword symbols. *IEEE Trans. Inform. Theory*, 58(11):6925–6934, Nov. 2012.
- [6] S. Hu, I. Tamo, and A. Barg. Combinatorial and lp bounds for lrc codes. In *Information Theory Proceedings (ISIT), 2016 IEEE International Symposium on*. IEEE, 2016.
- [7] E. Iceland and A. Samorodnitsky. On coset leader graphs of ldpc codes. *Information Theory, IEEE Transactions on*, 61(8):4158–4163, 2015.
- [8] R. McEliece, E. Rodemich, H. Rumsey, and L. Welch. New upper bounds on the rate of a code via the delarte-macwilliams inequalities. *Information Theory, IEEE Transactions on*, 23(2):157–166, 1977.
- [9] I. Tamo and A. Barg. A family of optimal locally recoverable codes. *IEEE Transactions on Information Theory*, 60(8):4661–4676, 2014.